arXiv:1607.01664v1 [stat.CO] 6 Jul 2016

Personalized Optimization for Computer Experiments with Environmental Inputs

Shifeng Xiong

Academy of Mathematics and Systems Science,

Chinese Academy of Sciences, Beijing 100190

Email: xiong@amss.ac.cn

Abstract Optimization problems with both control variables and environmental variables arise in many fields. This paper introduces a framework of personalized optimization to handle such problems. Unlike traditional robust optimization, personalized optimization devotes to finding a series of optimal control variables for different values of environmental variables. Therefore, the solution from personalized optimization consists of optimal surfaces defined on the domain of the environmental variables. When the environmental variables can be observed or measured, personalized optimization yields more reasonable and better solutions than robust optimization. The implementation of personalized optimization for complex computer models is discussed. Based on statistical modeling of computer experiments, we provide two algorithms to sequentially design input values for approximating the optimal surfaces. Numerical examples show the effectiveness of our algorithms.

KEY WORDS: Expected improvement; Kriging; Personalized medicine; Robust optimization; Sequential design.

1 Introduction

Optimization has broad applications in engineering, economics, operations research, and many other fields. A canonical form of optimization problems is

$$\min_{\mathbf{x}\in\mathcal{D}} f(\mathbf{x}), \tag{1}$$

where f is a real-valued function defined on a subset $\mathcal{D}$ of an Euclidean space. Since we use the minimization problem to represent general optimization problems in (1), the objective function f is called *cost* throughout this paper.

In many practical situations, the vector of variables $\mathbf{x}$ in (1) consist of two parts: control variables and environmental variables, denoted by $\mathbf{s}$ and $\mathbf{t}$, respectively. Control variables can be set as specific values by an engineer or scientist before the process that yields the value of f, while environmental variables, sometimes called noise variables, influence the cost in an uncertain way, depending on the user or on the environment at that time. There are a great number of examples of optimization with both control and environmental variables in quality engineering (Beyer and Sendhoff 2007) as well as in other fields (Ben-Tal, El Ghaoui, and Nemirovski 2009). A common case in medicine is as follows. Suppose that we will evaluate different medical treatments for a certain disease. These treatments can be represented by different values of control variables $\mathbf{s}$. For an individual patient, an index of the treatment effect corresponding to treatment $\mathbf{s}$ is denoted by f. Suppose that f also depends on some measurable information, denoted by the covaiates $\mathbf{t}$, of the patient. Here these covariates are actually environmental variables. It is desirable to find the best treatment, and this needs to optimize f with both $\mathbf{s}$ and $\mathbf{t}$.

Since environmental variables $\mathbf{t}$ cannot be controlled, the optimization problem in (1) should be modified to accommodate the situations where environmental variables exist. To this end, robust optimization techniques, which find the optimal *robust* value of $\mathbf{s}$ against the uncertainty from $\mathbf{t}$, have been well developed (Ben-Tal, El Ghaoui, and Nemirovski 2009). There are a number of definitions of this *robustness* with different focuses. In the above

medical example, the optimal robust solution can be defined as the treatment that makes the average effects over all patients under investigation the best.

A recent trend in medical research is *personalized medicine* (Jain 2015), also called precision medicine, stratified medicine, or P4 medicine. Roughly speaking, it means selection of treatment best suited for an individual according to the personal information of the individual, particularly the gene information. Personalized medicine has bred many new statistical issues like subgroup identification which a number of statisticians are working on (Ruberg, Chen, and Wang 2010; Xu et al. 2015; Zhao et al. 2015). Here we focus on optimization and let us return to the above example where the index f of the treatment effect depends on the covaiates $\mathbf{t}$. Under the spirit of personalized medicine, the best treatment should vary with the change of $\mathbf{t}$, and should optimize f for each given $\mathbf{t}$. It can be expected that the treatment from personalized medicine is generally better than that from robust optimization, which is the same for all $\mathbf{t}$. This idea can be borrowed to cope with general optimization problems with both control and environmental variables. This paper introduces a framework of *personalized optimization*, which devotes to finding a series of optimal control variables for different values of environmental variables. The solution from personalized optimization has a form of several curves or surfaces, called *profile optimal curves* (POC's) or *profile optimal surfaces* (POS's). Personalized optimization is applicable to the situations where the environmental variables can be observed or measured, and seems more reasonable than traditional robust optimization in practice.

When the cost function has an analytic expression, the problem of finding POS's (or POC's) can be solved for given environmental variables by traditional optimization techniques such as the Newton-Raphson algorithm (Lange 2010). In this paper we discuss this problem for computer experiments (Santner, Williams and Notz 2003). Nowadays computer experiments are commonly used to study a computer simulation in engineering and scientific investigations. Computer simulations usually have complex input-output relationships with long running times of the computer codes involved. A computer simulation can be viewed as an expensive black-box function when its output/response is scalar, A popular way to study such expensive function is to build a statistical surrogate model based on a small number

of output values on elaborately designed input values (Santner, Williams and Notz 2003). To study optimization algorithms for computer experiments, we need to integrate statistical methods with optimization techniques, and this becomes a hot area in both statistics and optimization because of the broad applications of computer experiments.

When only control inputs exist, Jones, Schonlau, and Welch (1998) proposed an expected improvement (EI) algorithm for optimizing expensive black-box functions. The EI method has proven to be a success, and its variants or related methods are proposed for various cases by many aothors; see Brochu, Cora, and De Freitas (2010), Picheny et al. (2013), and He, Tuo, and Wu (2016) among others. The EI method has also be extended to the situations where environmental inputs exist (Williams, Santner, and Notz 2000; Lehman, Santner, and Notz 2004; Marzat, Walter, and Piet-Lahanier 2013; Marzat, Walter, and Piet-Lahanier 2016). These extended EI methods are all within the framework of robust optimization, and search solutions that optimize certain robust criteria. Unlike them, in this paper we discuss the personalized optimization problem of computer experiments.

This paper is organized as follows. Section 2 introduces some notation and definitions in personalized optimization. Section 3 provides two sequential algorithms to implement personalized optimization for computer experiments. Section 4 presents numerical examples to evaluate the proposed algorithms. We end the paper with some discussion in Section 5.

2 Personalized optimization: A complement to robust optimization

Consider a cost function

$$y = f(\mathbf{s}, \mathbf{t}), \tag{2}$$

where $\mathbf{s} \in \mathcal{D}_{\mathrm{s}} \subset \mathbb{R}^p$ and $\mathbf{t} \in \mathcal{D}_{\mathrm{t}} \subset \mathbb{R}^q$ are control and environmental variables, respectively. There are two typical robust criteria in robust optimization. The first one is the expectation

criterion (Beyer and Sendhoff 2007)

$$f_{\mathrm{E}}(\mathbf{s}) = \int_{\mathcal{D}_{\mathrm{t}}} f(\mathbf{s}, \mathbf{t}) \varphi(\mathbf{t}) \, d\, \mathbf{t},$$

where $\varphi(\mathbf{t})$ is a probability density function on $\mathcal{D}_{\mathrm{t}}$. The robust solution according to this criterion is

$$\mathbf{s}_{\mathrm{E}}^{*} = \arg\min_{\mathbf{s} \in \mathcal{D}_{\mathrm{s}}} f_{\mathrm{E}}(\mathbf{s}). \tag{3}$$

A number of authors designed algorithms for approximating $\mathbf{s}_{\mathrm{E}}^{*}$ in computer experiments (Williams, Santner, and Notz 2000; Janusevskis and Le Riche 2013). The second robust criterion is the maximization criterion (also called worst-case criterion) (Ben-Tal and Nemirovski 1998)

$$f_{\mathrm{M}}(\mathbf{s}) = \max_{\mathbf{t} \in \mathcal{D}_{\mathrm{t}}} f(\mathbf{s}, \mathbf{t}).$$

Denote the corresponding solution by

$$\mathbf{s}_{\mathrm{M}}^{*} = \arg\min_{\mathbf{s} \in \mathcal{D}_{\mathrm{s}}} f_{\mathrm{M}}(\mathbf{s}). \tag{4}$$

Research articles on this criterion in computer experiments include Marzat, Walter, and Piet-Lahanier (2013) and Marzat, Walter, and Piet-Lahanier (2016), etc. There are other robust criteria used in practice; see Beyer and Sendhoff (2007).

When $\mathbf{t}$ can be observed or measured, the optimal $\mathbf{s}$ should minimize the cost for each given $\mathbf{t}$, and this is the main idea of personalized optimization. In other words, the purpose of personalized optimization is to find the *profile optimal surfaces* (POS's) (or called *profile optimal curves* (POC's) for $q = 1$)

$$\mathfrak{s}(\mathbf{t}) = \arg\min_{\mathbf{s} \in \mathcal{D}_{\mathrm{s}}} f(\mathbf{s}, \mathbf{t}). \tag{5}$$

The right side of (5) may be non-unique. For such cases, we assign any minimum to define $\mathfrak{s}(\mathbf{t})$. We call a mapping from $\mathcal{D}_{\mathrm{t}}$ to $\mathcal{D}_{\mathrm{s}}$ *personalized decision* in the following. Clearly, $\mathfrak{s}$ in (5) is a personalized decision. In addition, the robust solutions $\mathbf{s}_{\mathrm{E}}^{*}$ in (3) and $\mathbf{s}_{\mathrm{M}}^{*}$ in (4)

actually define two constant personalized decisions

$$u_{\mathrm{E}}(\mathbf{t}) \equiv \mathbf{s}_{\mathrm{E}}^* \quad \text{and} \quad u_{\mathrm{M}}(\mathbf{t}) \equiv \mathbf{s}_{\mathrm{M}}^*, \tag{6}$$

respectively. Since for any $\mathbf{s}$ and $\mathbf{t}$,

$$f(\mathfrak{s}(\mathbf{t}), \mathbf{t}) \leqslant f(\mathbf{s}, \mathbf{t}), \tag{7}$$

$\mathfrak{s}$ is the best personalized decision in the sense of minimizing the cost.

For a personalized decision u, define its expected cost and maximum cost as

$$C_{\mathrm{E}}(u) = \int_{\mathcal{D}_{\mathbf{t}}} f(u(\mathbf{t}), \mathbf{t}) \varphi(\mathbf{t})\, d\,\mathbf{t}, \tag{8}$$

$$C_{\mathrm{M}}(u) = \max_{\mathbf{t} \in \mathcal{D}_{\mathbf{t}}} f(u(\mathbf{t}), \mathbf{t}), \tag{9}$$

respectively. The expected cost and maximum cost can be used to evaluate the overall performance of u. We then use them to compare robust optimization and personalized optimization. By (7) and the definitions in (6), we have

$$C_{\mathrm{E}}(\mathfrak{s}) \leqslant C_{\mathrm{E}}(u_{\mathrm{E}}) \leqslant C_{\mathrm{E}}(u_{\mathrm{M}})$$

and

$$C_{\mathrm{M}}(\mathfrak{s}) \leqslant C_{\mathrm{M}}(u_{\mathrm{M}}) \leqslant C_{\mathrm{M}}(u_{\mathrm{E}}).$$

Consider a simple example with $f(s,t) = (s-t)^2,\ s,t \in [0,1]$. Some algebra gives that $C_{\mathrm{E}}(\mathfrak{s}) = C_{\mathrm{M}}(\mathfrak{s}) = 0$, $C_{\mathrm{E}}(u_{\mathrm{E}}) = 1/12$, and $C_{\mathrm{M}}(u_{\mathrm{M}}) = 1/4$.

3 Sequential algorithms for searching POS's of computer models

This section discusses personalized optimization for computer experiments. Since computer models are often expensive black-box functions, we first describe the popular statistical modeling methods for them in Section 3.1. In Section 3.2 we present two sequential algorithms to search POS's of computer models.

3.1 The Gaussian process model

The Gaussian process model (also called Kriging model) (Matheron 1963) is widely used to analyze computer experiments (Sacks et al. 1989). It models the output of a computer simulation as

$$f(\mathbf{x}) = \mathbf{g}(\mathbf{x})'\boldsymbol{\beta} + Z(\mathbf{x}), \tag{10}$$

where $\mathbf{x} \in \mathcal{D} \subset \mathbb{R}^d$, $\mathbf{g}(\mathbf{x}) = (g_1(\mathbf{x}), \ldots, g_m(\mathbf{x}))'$ is a pre-specified set of functions, $\boldsymbol{\beta}$ is a vector of unknown regression coefficients, and $Z(\mathbf{x})$ is a stationary Gaussian process with mean zero, variance σ^2, and correlation parameters $\boldsymbol{\theta}$, denoted by $\mathrm{GP}(0, \sigma^2, \boldsymbol{\theta})$. The covariance between $Z(\mathbf{x}_1)$ and $Z(\mathbf{x}_2)$ in (10) is represented by

$$\mathrm{Cov}[Z(\mathbf{x}_1), Z(\mathbf{x}_2)] = \sigma^2 R(\mathbf{x}_1 - \mathbf{x}_2 \,|\, \boldsymbol{\theta}), \tag{11}$$

where $R(\cdot \,|\, \boldsymbol{\theta})$ is the correlation function depending on a parameter vector $\boldsymbol{\theta}$. A popular choice of R is the Gaussian correlation function

$$R(\mathbf{u} \,|\, \theta) = \exp(-\theta_i u_i^2) \quad \text{for } \mathbf{u} = (u_1, \ldots, u_d)' \in \mathcal{D}, \tag{12}$$

where $\theta_1, \ldots, \theta_d > 0$ are the correlation parameters.

The parameters in model (10) can be estimated by the maximum likelihood method. Suppose the set of input values is $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$, where $\mathbf{x}_i = (x_{i1}, \ldots, x_{id})'$ for $i = 1, \ldots, n$. The corresponding response values $\mathbf{y} = (f(\mathbf{x}_1), \ldots, f(\mathbf{x}_n))'$. The negative log-likelihood, up to an additive constant, is proportional to

$$n \log(\sigma^2) + \log(\det(\mathbf{R})) + (\mathbf{y} - \mathbf{G}\boldsymbol{\beta})'\mathbf{R}^{-1}(\mathbf{y} - \mathbf{G}\boldsymbol{\beta})/\sigma^2, \tag{13}$$

where $\mathbf{R}$ is the $n \times n$ correlation matrix whose (i,j)th entry is $R(\mathbf{x}_i - \mathbf{x}_j \,|\, \boldsymbol{\theta})$ defined in (11), "det" denotes matrix determinant, and $\mathbf{G} = (\mathbf{g}(\mathbf{x}_1), \ldots, \mathbf{g}(\mathbf{x}_n))'$.

When $\boldsymbol{\theta}$ is known, the maximum likelihood estimators (MLEs) of $\boldsymbol{\beta}$ and σ^2 are

$$\begin{cases} \hat{\boldsymbol{\beta}} = (\mathbf{G}'\mathbf{R}^{-1}\mathbf{G})^{-1}\mathbf{G}'\mathbf{R}^{-1}\mathbf{y}, \\ \hat{\sigma}^2 = (\mathbf{y} - \mathbf{G}\hat{\boldsymbol{\beta}})'\mathbf{R}^{-1}(\mathbf{y} - \mathbf{G}\hat{\boldsymbol{\beta}})/n. \end{cases} \tag{14}$$

For an untried point $\mathbf{x}_0 \in \mathcal{D}$, the best linear unbiased predictor $\hat{f}$ of f (Santner, Williams and Notz 2003) is

$$\hat{f}(\mathbf{x}_0) = \mathbf{g}(\mathbf{x}_0)'\hat{\boldsymbol{\beta}} + \mathbf{r}_0'\mathbf{R}^{-1}\big(\mathbf{y} - \mathbf{G}\hat{\boldsymbol{\beta}}\big), \tag{15}$$

where $\mathbf{r}_0 = \big(R(\mathbf{x}_0 - \mathbf{x}_1 \,|\, \boldsymbol{\theta}), \ldots, R(\mathbf{x}_0 - \mathbf{x}_n \,|\, \boldsymbol{\theta})\big)'$. Given $\alpha \in (0,1)$, the $(1-\alpha)$ prediction interval of $f(\mathbf{x}_0)$ is

$$P\left(f(\mathbf{x}_0) \in \hat{f}(\mathbf{x}_0) \pm \phi(\mathbf{x}_0)\, t_{n-d}(\alpha/2)\right) = 1 - \alpha, \tag{16}$$

where $\phi(\mathbf{x}_0) \geqslant 0$,

$$\begin{aligned} \phi(\mathbf{x}_0)^2 &= \frac{Q^2}{n-d}\left\{1 - (\mathbf{g}(\mathbf{x}_0)', \mathbf{r}_0')\begin{pmatrix} \mathbf{0} & \mathbf{G}' \\ \mathbf{G} & \mathbf{R} \end{pmatrix}^{-1}\begin{pmatrix} \mathbf{g}(\mathbf{x}_0) \\ \mathbf{r}_0 \end{pmatrix}\right\}, \\ Q^2 &= \mathbf{y}'\big[\mathbf{R}^{-1} - \mathbf{R}^{-1}\mathbf{G}(\mathbf{G}'\mathbf{R}^{-1}\mathbf{G})^{-1}\mathbf{G}'\mathbf{R}^{-1}\big]\mathbf{y}, \end{aligned} \tag{17}$$

and $t_{n-d}(\alpha/2)$ is the upper $\alpha/2$ quantile of Student's t-distribution with $n-d$ degrees of freedom. Denote the lower bound in (16) by

$$L(\mathbf{x}_0) = \hat{f}(\mathbf{x}_0) - \phi(\mathbf{x}_0)\, t_{n-d}(\alpha/2). \tag{18}$$

When $\boldsymbol{\theta}$ is unknown, by plugging (14) into (13), we have the MLE of $\boldsymbol{\theta}$

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}}\, n\log(\hat{\sigma}^2) + \log\left(\det(\mathbf{R})\right).$$

The predictor $\hat{f}$ in (15) and the prediction intervals in (16) can be modified by replacing $\boldsymbol{\theta}$ with $\hat{\boldsymbol{\theta}}$.

Let us return to the computer model having the form in (2), where $\mathbf{x}$ in (10) is $(\mathbf{s}', \mathbf{t}')'$, $d = p+q$, and $\mathcal{D} = \mathcal{D}_{\mathrm{s}} \times \mathcal{D}_{\mathrm{t}}$. Suppose that we have n pairs of input values $(\mathbf{s}_1', \mathbf{t}_1')', \ldots, (\mathbf{s}_n', \mathbf{t}_n')'$, and the corresponding output values $\mathbf{y} = (f(\mathbf{s}_1, \mathbf{t}_1), \ldots, f(\mathbf{s}_n, \mathbf{t}_n))'$. By (15), the predictor $\hat{f}$ can be obtained based on the data. Consequently, the POS's can be estimated by

$$\hat{\mathfrak{s}}(\mathbf{t}) = \arg\min_{\mathbf{s} \in \mathcal{D}_{\mathrm{s}}} \hat{f}(\mathbf{s}, \mathbf{t}). \tag{19}$$

3.2 Algorithm description

This subsection presents two sequential algorithms to search POS's. Since more experimental points closed to the true POS's can make their estimators in (19) more accurate, the basic idea of our algorithms is to sequentially select the next point according to a criterion that measures the proximity of a point to the POS's.

First we give two useful definitions. For $\mathbf{t} \in \mathcal{D}_{\mathrm{t}}$, define

$$\tilde{\mathfrak{s}}(\mathbf{t}) = \arg\min_{\mathbf{s} \in \mathcal{D}_{\mathrm{s}}} L(\mathbf{s}, \mathbf{t}), \tag{20}$$

where L is defined in (18). For a set $\mathcal{X}$ of points, define

$$M(\mathbf{x}, \mathcal{X}) = \min_{\mathbf{a} \in \mathcal{X}} \|\mathbf{x} - \mathbf{a}\|. \tag{21}$$

The proposed algorithms are hierarchical. Suppose that the current experimental points are $(\mathbf{s}_1', \mathbf{t}_1')', \ldots, (\mathbf{s}_n', \mathbf{t}_n')'$. For given $\mathbf{t}_{n+1}$, the goal is to optimize f with respect to $\mathbf{s}$. Hence the idea of EI (Jones, Schonlau, and Welch 1998) can be used. Unfortunately, since there is no, or only a part of, experimental points on the profile $\{(\mathbf{s}', \mathbf{t}')' \in \mathcal{D} : \ \mathbf{t} = \mathbf{t}_{n+1}\}$, the EI criterion is not well defined for our problem. Here we adopt the confidence bound criterion (Cox and John 1997) in (20) to find the next $\mathbf{s}$, that is, to use $\mathbf{s}_{n+1}$ that minimizes the lower bound of the $(1-\alpha)$ prediction interval of $f(\mathbf{s}, \mathbf{t}_{n+1})$. The confidence bound criterion takes

the uncertainty in predicting the response on untried points into account like EI, and has a simple form for our problem. The nominal level α controls the selection between the global solution and local solution. For smaller value of α, greater weight in the confidence bound criterion will be placed on the variance of the predictor, and thus we tend to approximate the global minimizer. Such a search, however, can be slow to reach a local solution. Therefore, smaller α is more suitable for multimodal functions.

We next discuss the selection of $\mathbf{t}_{n+1}$. Two selection methods are presented, which form the difference between the two proposed algorithms SHA1 and SHA2; see Algorithms 1 and 2. The first method is model free through optimizing a space-filling criterion in (21). This criterion stems from the maximin distance criterion (Johnson, Moore, and Ylvisaker 1990). The second method is model-based; see Step 3 in SHA2. It selects the next $\mathbf{t}$ with the maximal amount of uncertainty at the point which consists of $\mathbf{t}$ and $\mathbf{s}$ selected from the confidence bound criterion. We will compare SHA1 and SHA2 with different α via numerical examples in the next section.

Algorithm 1 The sequential hierarchical algorithm 1 (SHA1)

Inputs:

The initial sample size n_0 and the coverage level $\alpha \in (0, 1)$.

Steps:

1. Choose the initial design $\mathcal{X} = \{(\mathbf{s}_1', \mathbf{t}_1')', \ldots, (\mathbf{s}_{n_0}', \mathbf{t}_{n_0}')'\} \subset \mathcal{D}_\mathrm{s} \times \mathcal{D}_\mathrm{t}$ and compute the corresponding responses $\mathbf{y} = (y_1, \ldots, y_{n_0})'$. Let $n = n_0$.

For $n = n_0, n_0 + 1, \ldots,$

2. Fit the Gaussian process model with the current data $\{\mathcal{X}, \mathbf{y}\}$.
3. Find
$$\mathbf{t}_{n+1} = \arg\max_{\mathbf{t} \in \mathcal{D}_\mathrm{t}} M\big(\mathbf{t}, \mathcal{T}_n\big),$$
where $\mathcal{T}_n = \{\mathbf{t}_1, \ldots, \mathbf{t}_n\}$ and M is defined in (21). Let $\mathbf{s}_{n+1} = \tilde{\mathfrak{s}}(\mathbf{t}_{n+1})$.
4. Compute $y_{n+1} = f(\mathbf{s}_{n+1}, \mathbf{t}_{n+1})$.
5. Let $\mathcal{X} \leftarrow \mathcal{X} \cup \{(\mathbf{s}_{n+1}', \mathbf{t}_{n+1}')'\}$, $\mathbf{y} \leftarrow (\mathbf{y}', y_{n+1})'$, and $n \leftarrow n + 1$.

End for if some stopping rule is met.

In SHA1 and SHA2, the initial design can be chosen as a space-filling design (Joseph 2016) which are commonly used in computer experiments. Such designs include the maximin distance design (Johnson, Moore, and Ylvisaker 1990), Latin hypercube design (Mckay,

Algorithm 2 The sequential hierarchical algorithm 2 (SHA2)

Inputs: The same as in Algorithm 1.

Steps:

All steps are the same as in Algorithm 1 except Step 3: Find

$$\mathbf{t}_{n+1} = \arg\max_{\mathbf{t}\in\mathcal{D}_{\mathrm{t}}} \phi\big(\tilde{\mathfrak{s}}(\mathbf{t}), \mathbf{t}\big),$$

where ϕ and $\tilde{\mathfrak{s}}$ are defined in (17) and (20), respectively. Let $\mathbf{s}_{n+1} = \tilde{\mathfrak{s}}(\mathbf{t}_{n+1})$.

Beckman, and Conover 1979), and uniform design (Fang et al. 2000). We use the Gaussian correlation function (12) to build the Gaussian process model in (10) throughout, where $\mathbf{g}(\mathbf{x})$ in (10) is set as $(1, s_1, \ldots, s_p, t_1, \ldots, t_q)'$. Let $\hat{f}_n$ denote the predictor by (15) based on the current data in the nth iteration, and then the estimators of POS's, $\hat{\mathfrak{s}}_n(\mathbf{t})$, can be computed by (19). We stop the iterations in SHA1 and SHA2 if

$$\int_{\mathcal{D}_{\mathrm{t}}} \frac{|\hat{f}_{n-1}(\hat{\mathfrak{s}}_{n-1}(\mathbf{t}), \mathbf{t}) - \hat{f}_n(\hat{\mathfrak{s}}_n(\mathbf{t}), \mathbf{t})|}{|\hat{f}_n(\hat{\mathfrak{s}}_n(\mathbf{t}), \mathbf{t})|} \varphi(\mathbf{t})\, d\,\mathbf{t} < \varepsilon_1 \quad \text{and} \quad \int_{\mathcal{D}_{\mathrm{t}}} \phi(\hat{\mathfrak{s}}_n(\mathbf{t}), \mathbf{t}) \varphi(\mathbf{t})\, d\,\mathbf{t} < \varepsilon_2,$$

where ε_1 and ε_2 are pre-specified thresholdings. The stopping rule can also be set as

$$\max_{\mathbf{t}\in\mathcal{D}_{\mathrm{t}}} \frac{|\hat{f}_{n-1}(\hat{\mathfrak{s}}_{n-1}(\mathbf{t}), \mathbf{t}) - \hat{f}_n(\hat{\mathfrak{s}}_n(\mathbf{t}), \mathbf{t})|}{|\hat{f}_n(\hat{\mathfrak{s}}_n(\mathbf{t}), \mathbf{t})|} < \varepsilon_1 \quad \text{and} \quad \max_{\mathbf{t}\in\mathcal{D}_{\mathrm{t}}} \phi(\hat{\mathfrak{s}}_n(\mathbf{t}), \mathbf{t}) < \varepsilon_2.$$

Besides, a practical and simple stopping rule in SHA1 and SHA2 is to stop the iterations until the run size reaches our budget.

4 Numerical illustrations

4.1 A simple two-dimensional example

Consider the simple function $f(s,t) = (s-t)^2,\ s,t \in [0,1]$, which is used in Section 2. It is obvious that the POC of f is $s = t,\ t \in [0,1]$. We begin with seven randomly generated points, and use SHA1 and SHA2 to approximate the POC. Figure 1 shows the locations of the added runs after seven iterations. We can see that these points are closed to the POC,

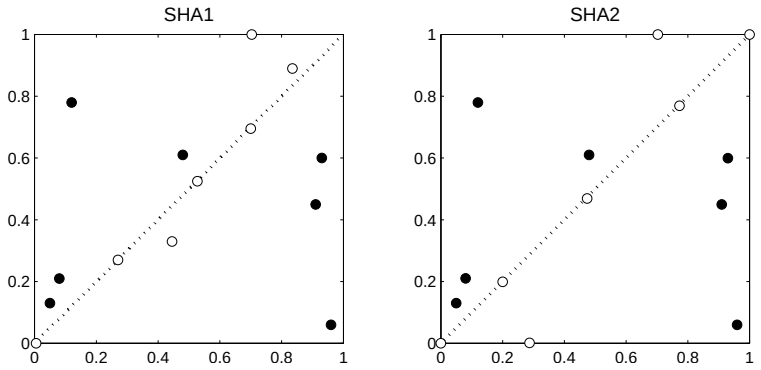

Figure 1: The designs corresponding to SHA1 and SHA2 ($\alpha = 0.2$) in Section 4.1. The black dots constitute the initial design, and the white dots denote the added runs after seven iterations. The dotted line is the true POC.

and this indicates that our algorithms work well.

4.2 Comparison of the costs from different methods

Let $\mathcal{D}_{\mathrm{s}} = [0,1]^p$, $\mathcal{D}_{\mathrm{t}} = [0,1]^q$, and φ in (8) be the density of the uniform distribution on $[0,1]^q$. We use the following six test functions:

$$
\begin{aligned}
f_1(s,t) &= 2|s^3 - t| + \exp(t)(s - 2t)^2, \\
f_2(s,t) &= \cos(10\sqrt{s^2+t^2})/(\sqrt{s^2+t^2}+1), \\
f_3(s,t) &= \min\{3 - 2s + 3t,\ 3 + 2s - t\}, \\
f_4(s,t) &= \left\{15t - \frac{5.1}{4\pi^2}(15s-5)^2 + \frac{5}{\pi}(15s-5) - 6\right\}^2 + 10\left(1 - \frac{1}{8\pi}\right)\cos(15s-5) + 10, \\
f_5(\mathbf{s},\mathbf{t}) &= (s_1 - |t_1 - t_2|)^2 + \left(s_2 - \sqrt{(t_1^2+t_2^2)/2}\right)^4, \quad (p=q=2) \\
f_6(\mathbf{s},\mathbf{t}) &= \sin(5s_1^2)(t_1 + 2s_2) - \cos(5s_3^2)/\sqrt{1+s_4^2} - 2t_2(s_1 - s_4), \quad (p=4, q=2)
\end{aligned}
$$

some of which have been used in the literature; see Jones, Schonlau, and Welch (1998) and Marzat, Walter, and Piet-Lahanier (2016), among others. We compare three sequential design strategies, SHA1, SHA2, and the Sobol' sequence (Niederreiter 1992) in approximation to POS's. The Sobol' sequence can be viewed as a sequential uniform design, which places

the points more and more uniformly over the experimental region as the number of points increases. In SHA1 and SHA2, the initial design is generated by the first n_0 points of the Sobol' sequence, where $n_0 = 10$ for $f_1, \ldots, f_4$ and $n_0 = 20$ for f_5 and f_6. Three nominal levels, $\alpha = 0.2,\ 0.5$, and 0.8, are used in SHA1 and SHA2. The expected costs in (8) and the maximum costs in (9) of the estimated POS's, $\hat{\mathfrak{s}}$ in (19), based on the three design strategies are computed in each iteration for the comparison. To compare personalized optimization and robust optimization, we also compute the expected and maximum costs of the constant personalized decisions $u_{\mathrm{E}}(\mathbf{t})$ and $u_{\mathrm{M}}(\mathbf{t})$ in (6) corresponding to the solutions from robust optimization. Here $u_{\mathrm{E}}(\mathbf{t})$ and $u_{\mathrm{M}}(\mathbf{t})$ are accurately computed by the forms of the cost functions. Therefore, they are better than their data-based estimators.

The numerical results for 30 iterations are presented in Figures 2 – 7. We can see that, the three methods for approximating POS's are better than the constant personalized decisions in most cases. In addition, SHA1 and SHA2 with larger α are generally better than the Sobol' sequence. For the two proposed algorithms, SHA2 seems to have slightly better overall performance than SHA1. The numerical results indicate that, first, the methods of personalized optimization, even not accurate, are generally better than the accurate solutions of robust optimization; second, the sequential algorithms designed for searching POS's can approximate the POS's better than space-filling designs.

5 Discussion

This paper has introduced the framework of personalized optimization to handle optimization problems with both control variables and environmental variables. We have discussed its implementation for computer experiments. Two sequential algorithms for searching POS's of complex computer models have been proposed. Theoretical and numerical analysis indicates that, when the environmental variables can be observed or measured, personalized optimization yields more reasonable and better solutions than traditional robust optimization.

Personalized optimization for computer experiments can be studied further in several

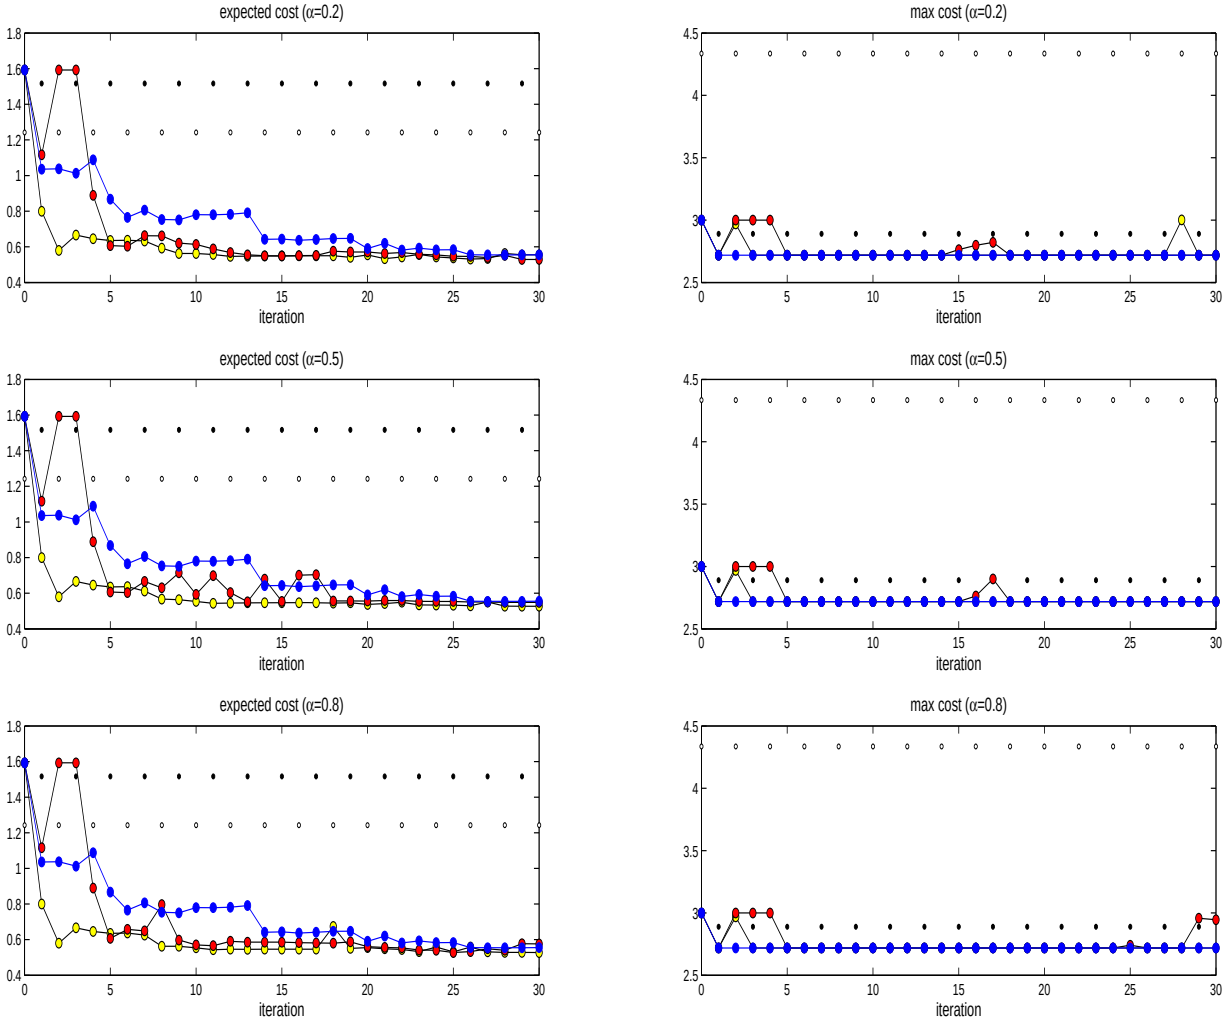

Figure 2: The costs concerning f_1 in Section 4.2. Legend: white dots — u_{E}, black dots — u_{M}, blue dots — Sobol' sequence, red dots — SHA1, yellow dots — SHA2.

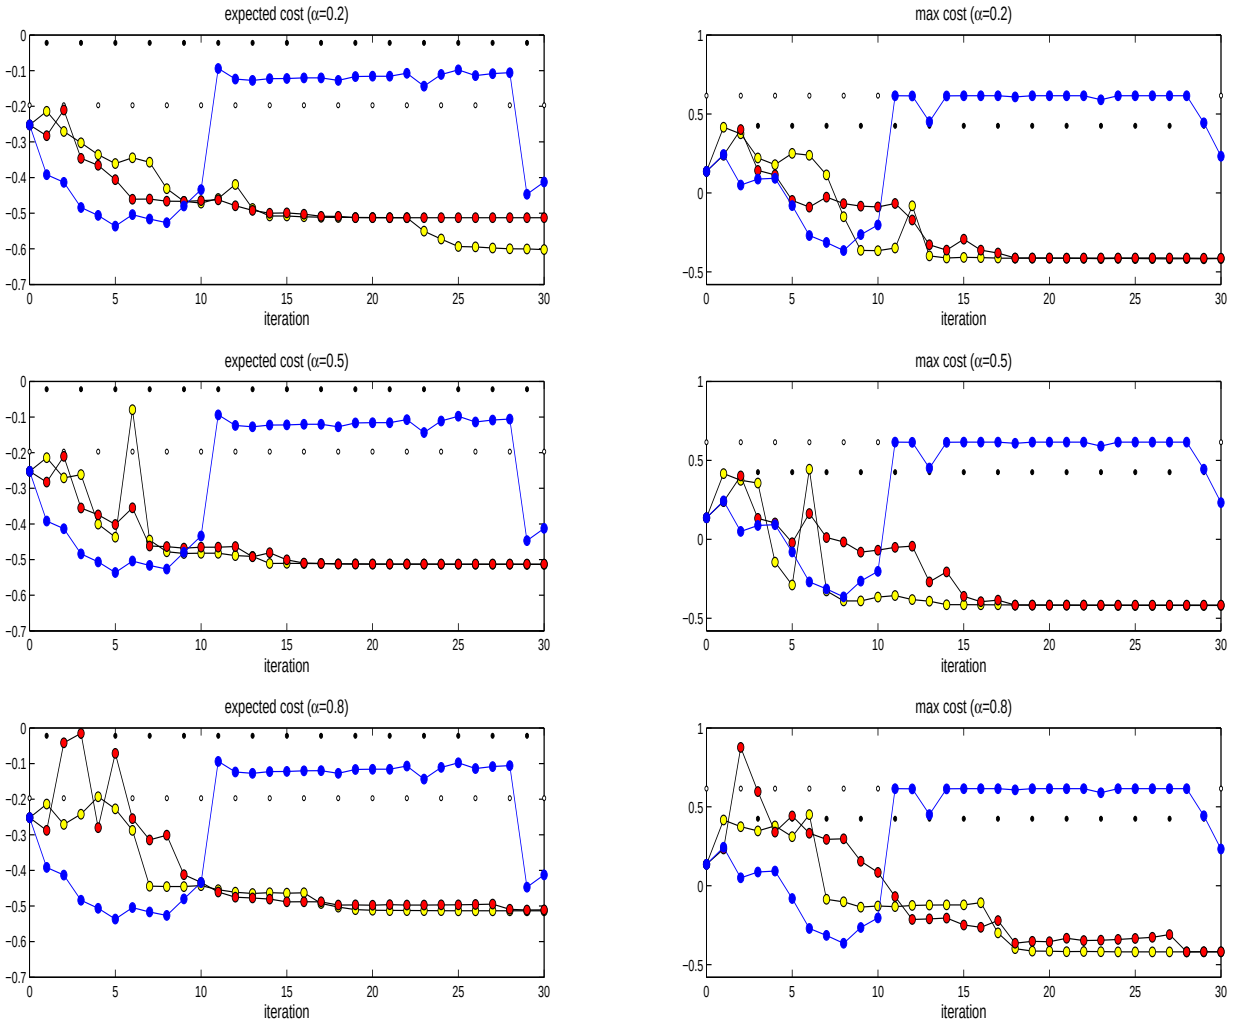

Figure 3: The costs concerning f_2 in Section 4.2. Legend: white dots — u_{E}, black dots — u_{M}, blue dots — Sobol' sequence, red dots — SHA1, yellow dots — SHA2.

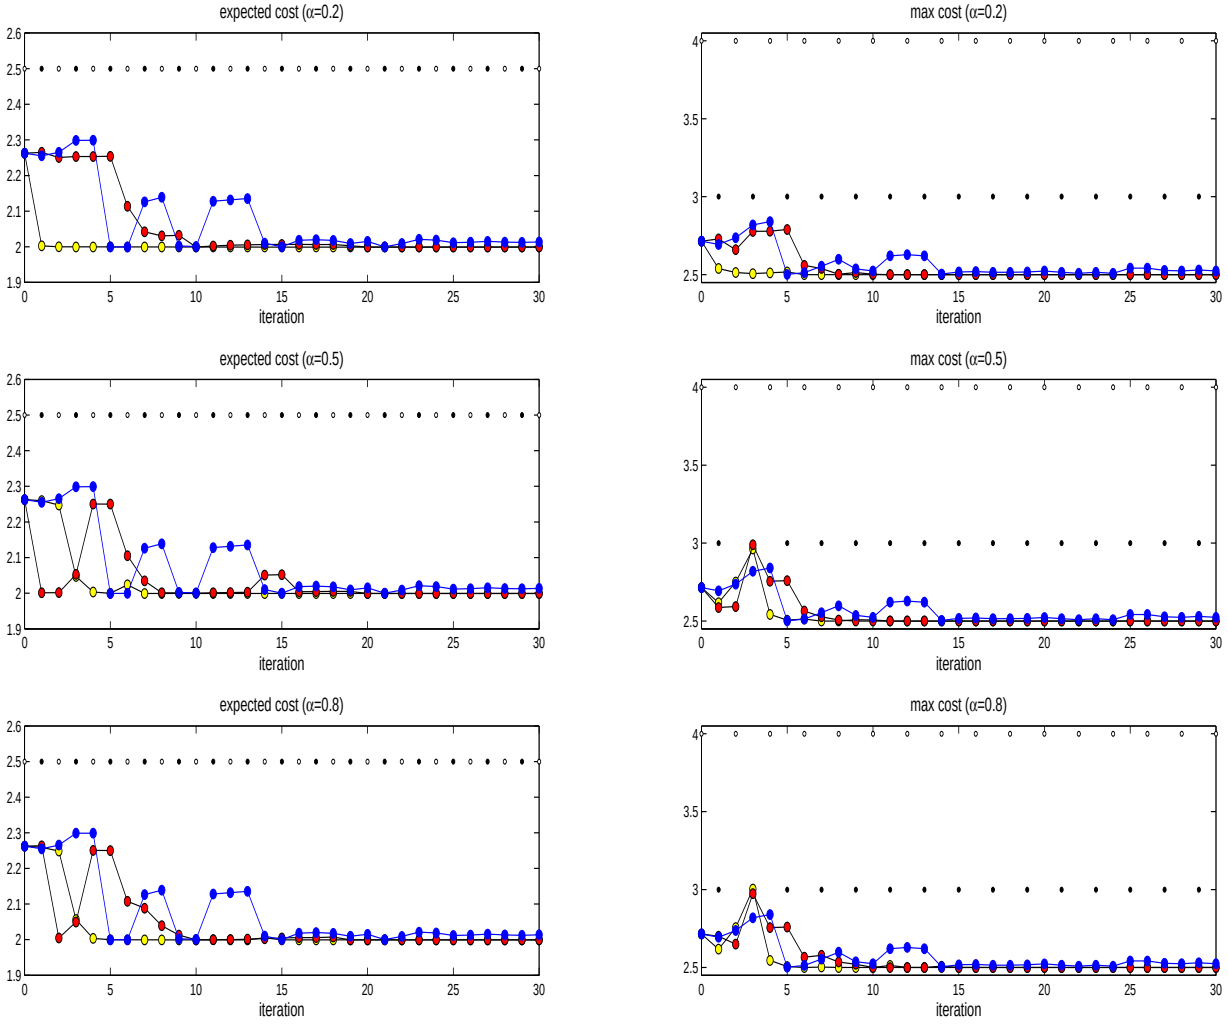

Figure 4: The costs concerning f_3 in Section 4.2. Legend: white dots — u_{E}, black dots — u_{M}, blue dots — Sobol' sequence, red dots — SHA1, yellow dots — SHA2.

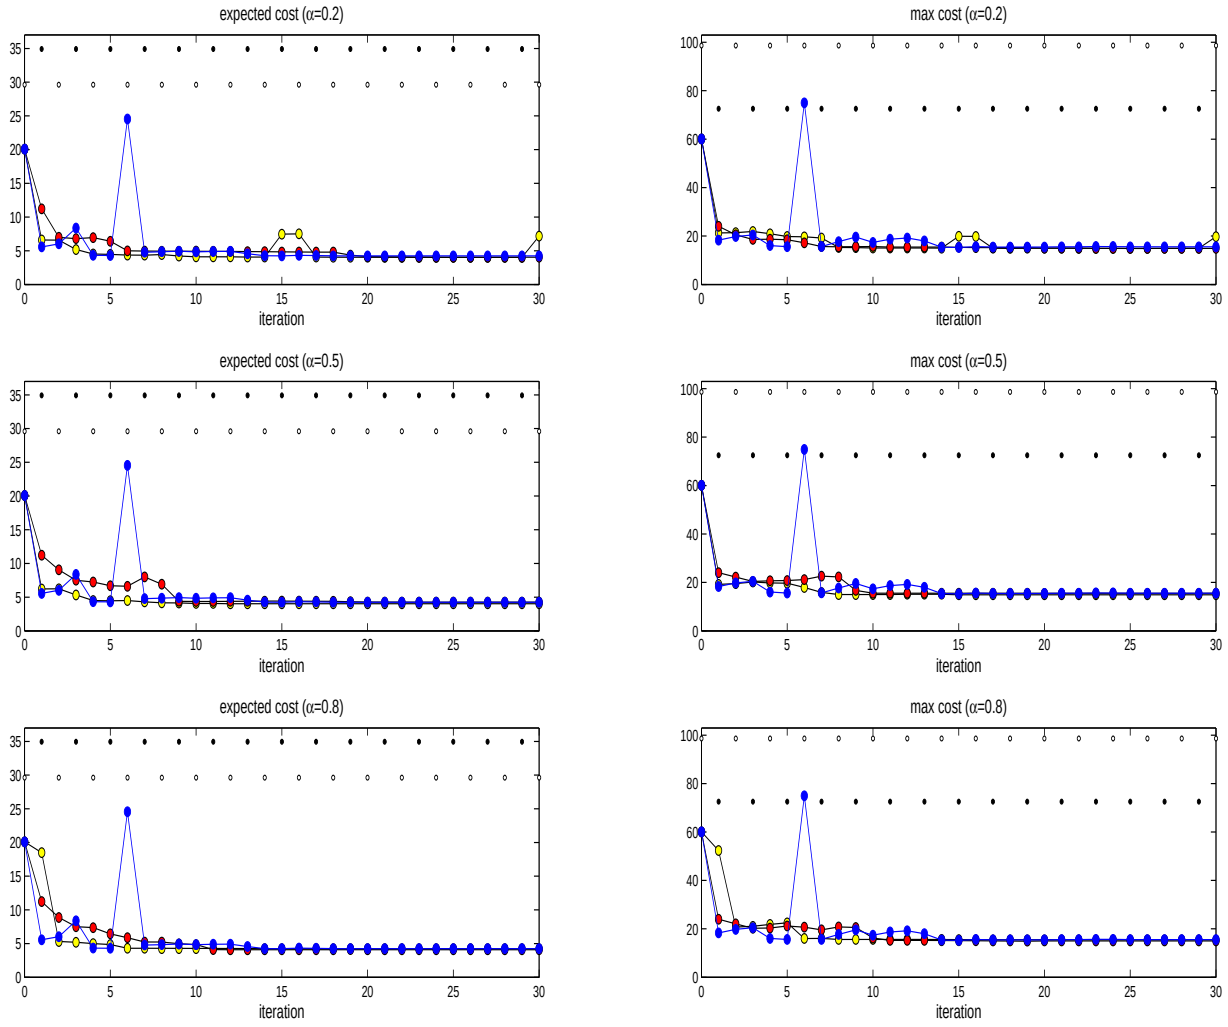

Figure 5: The costs concerning f_4 in Section 4.2. Legend: white dots — u_{E}, black dots — u_{M}, blue dots — Sobol' sequence, red dots — SHA1, yellow dots — SHA2.

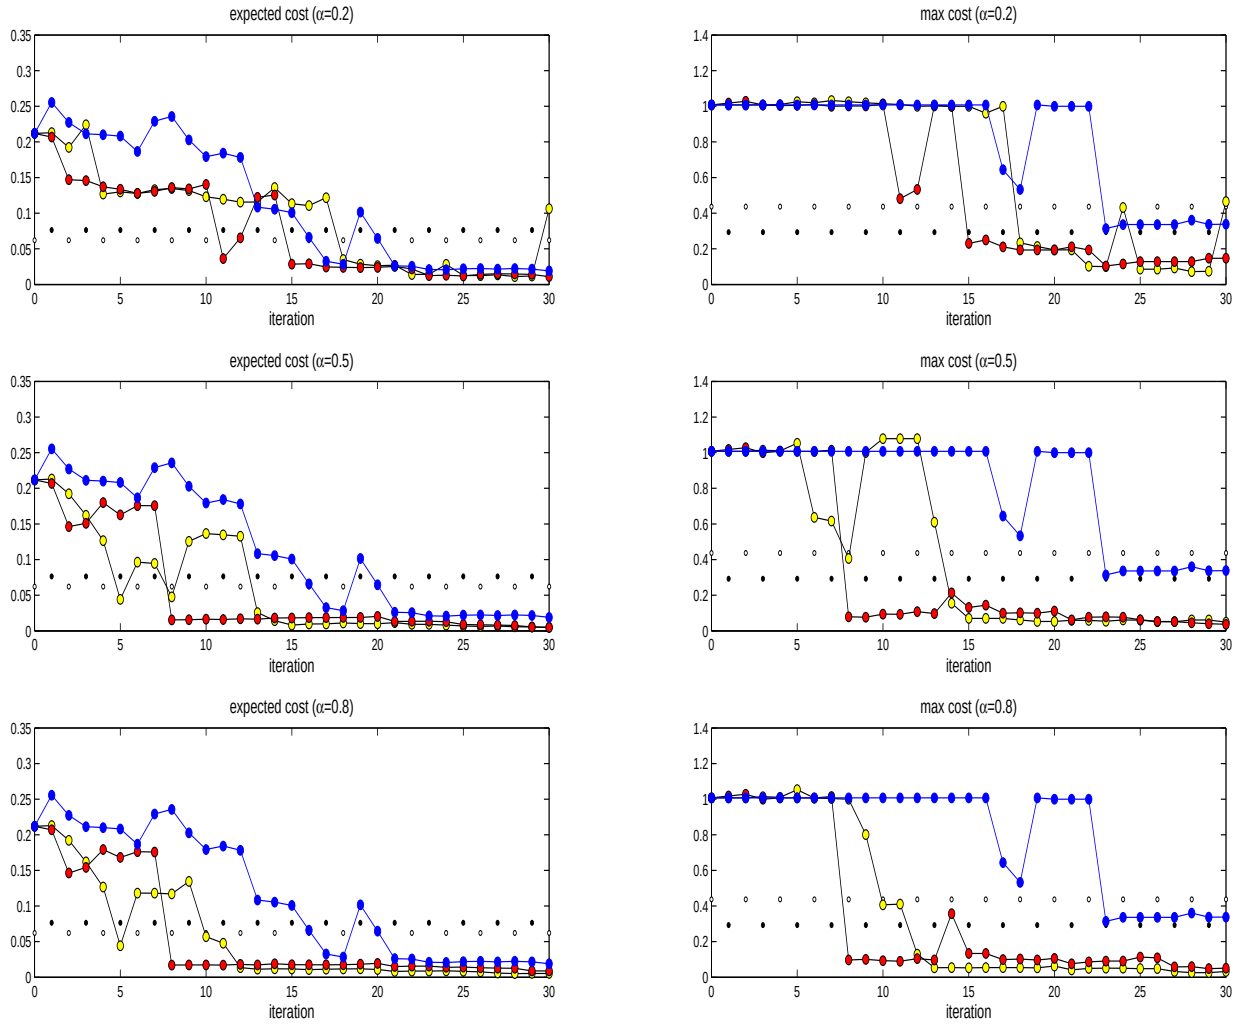

Figure 6: The costs concerning f_5 in Section 4.2. Legend: white dots — u_{E}, black dots — u_{M}, blue dots — Sobol' sequence, red dots — SHA1, yellow dots — SHA2.

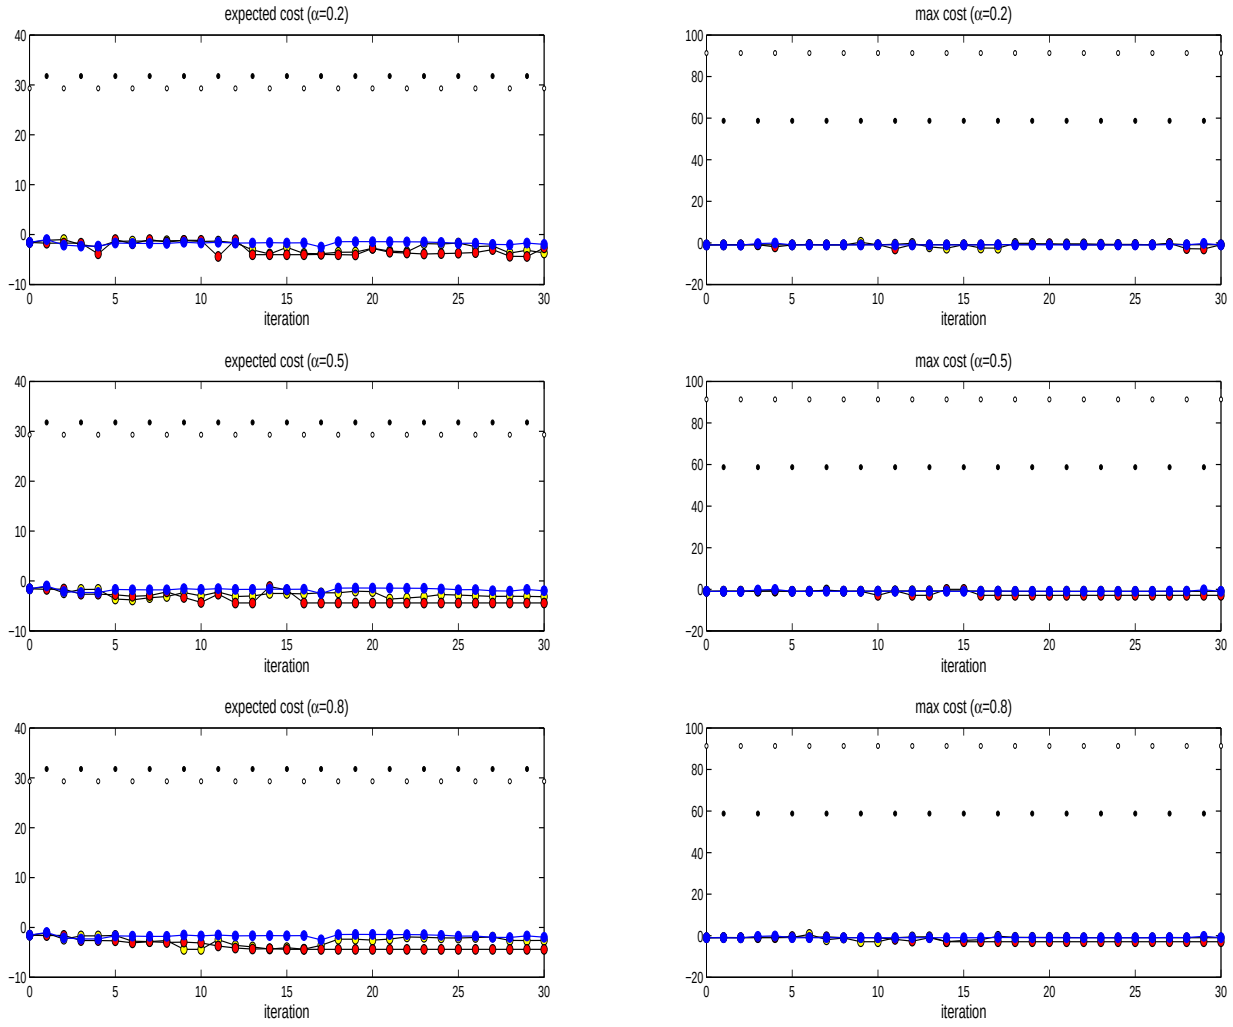

Figure 7: The costs concerning f_6 in Section 4.2. Legend: white dots — u_{E}, black dots — u_{M}, blue dots — Sobol' sequence, red dots — SHA1, yellow dots — SHA2.

directions. First, for large p and/or q, it is difficult to solve the optimization problems that are needed in personalized optimization such as those in (19) and Step 3 in SHA1. This is a common difficulty in optimization. Some global optimization methods can be tried, including multiple-startpoint methods (Ugray et al. 2007) and random search algorithms (Kirkpatrick, Gelatt, and Vecchi 1983; Dorsey and Mayer 1995). Sensitivity analysis techniques (Saltelli et al. 2008) can also be used to reduce the dimensionality of the initial optimization problems. Second, we can use different nominal level α for different $\mathbf{t}$ in SHA1 and SHA2. This modification may be better when the cost function has different levels of multimodality for different $\mathbf{t}$. Third, the proposed algorithms can be extended to the cases where both measurable and unmeasurable environmental variables exist. We need to combine our methods with robust optimization techniques for such cases. Fourth, since SHA1 and SHA2 are proposed to search optimal surfaces, there is not straightforward mathematical tool to study their convergence. It is valuable to develop a theoretical framework to fill this gap in the future.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant No. 11271355, 11471172) and Key Laboratory of Systems and Control, Chinese Academy of Sciences.

REFERENCES

Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. (2009). *Robust Optimization*. Princeton University Press: Princeton and Oxford.

Ben-Tal, A. and Nemirovski, A. (1998). Robust convex optimization. *Mathematics of Operations Research*, 23, 769–805.

Beyer, H. G. and Sendhoff, B. (2007). Robust optimization: a comprehensive survey. *Computer methods in applied mechanics and engineering*, 196, 3190–3218.

Brochu, E., Cora, V. M., and De Freitas, N. (2010). A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. arXiv preprint arXiv:1012.2599.

Cox, D. D. and John, S., (1997). SDO: A statistical method for global optimization. In M. N. Alexandrov and M. Y. Hussaini, editors, *Multidisciplinary Design Optimization: State of the Art*, pages 315–329. SIAM.

Dorsey, R. E. and Mayer, M. J. (1995). Genetic algorithms for estimation problems with multiple optima, nondifferentiability, and other irregular features. *Journal of Business and Economic Statistics*, 13, 53–66.

Fang, K. T., Lin, D. K. J., Winker, P.. and Zhang, Y. (2000). Uniform design: theory and application. *Technometrics*, 42, 953–963.

He, X., Tuo, R., and Wu, C. F. J. (2016). Optimization of multi-fidelity computer experiments via the EQIE criterion, *Technometrics*, In press.

Jain, K. K. (2015). *Textbook of Personalized Medicine*, second edition. New York: Springer.

Janusevskis, J. and Le Riche, R. (2013): Simultaneous Kriging-based estimation and optimization of mean response. *Journal of Global optimization*, 55, 313–336.

Johnson, M. E., Moore, L. M., and Ylvisaker, D. (1990). Minimax and maximin distance designs. *Journal of Statistical Planning and Inference*, 26, 131–148.

Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13, 455-492.

Joseph, V. R. (2016). Space-filling designs for computer experiments: a review, *Quality Engineering*, 28, 28–35.

Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220, 671–680.

Lange, K. (2010). *Numerical Analysis for Statisticians*. Springer Science & Business Media.

Lehman, J. S., Santner, T. J., and Notz, W. I. (2004). Designing computer experiments to determine robust control variables. *Statistica Sinica*, 14, 571–590.

Marzat, J., Walter, E., and Piet-Lahanier, H. (2013). Worst-case global optimization of black-box functions through Kriging and relaxation. *Journal of Global Optimization*, 55, 707–727.

Marzat, J., Walter, E., and Piet-Lahanier, H. (2016). A new expected-improvement algorithm for continuous minimax optimization. *Journal of Global Optimization*, In press.

McKay, M. D., Beckman, R. J. and Conover, W. J. (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21, 239–245.

Niederreiter, H., (1992). *Random Number Generation and Quasi-Monte Carlo Methods*, NSFCBMS Regional Conference on Random Number Generation. (Vol. 63). Philadelphia: Society for Industrial and Applied mathematics.

Picheny, V., Ginsbourger, D., Richet, Y., and Caplin, G. (2013), Quantile-based optimization of noisy computer experiments with tunable precision, *Technometrics*, 55, 2–13.

Ruberg, S. J., Chen, L., and Wang, Y. (2010). The mean does not mean as much anymore: Finding sub-groups for tailored therapeutics. *Clinical Trials*, 7, 574–583.

Sacks, J., Welch, W. J., Mitchell, T. J. and Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical Science*, 4, 409–423.

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M. and Tarantola, S. (2008). *Global Sensitivity Analysis: The Primer*. John Wiley & Sons.

Santner, T. J., Williams, B. J. and Notz, W. I. (2003). *The Design and Analysis of Computer Experiments*. Springer, New York.

Ugray, Z., Lasdon, L., Plummer, J., Glover, F., Kelly, J., and Marti, R. (2007). Scatter search and local NLP solvers: a multistart framework for global optimization. *INFORMS Journal on Computing*, 19, 328–340.

Williams, B. J., Santner, T. J., and Notz, W. I. (2000). Sequential design of computer experiments to minimize integrated response functions. *Statistica Sinica*, 10, 1133–1152.

Xu, Y., Yu, M., Zhao, Y. Q., Li, Q., Wang, S., and Shao, J. (2015). Regularized outcome weighted subgroup identification for differential treatment effects. *Biometrics*, 71, 645–653.

Zhao, Y. Q., Zeng, D., Laber, E. B., and Kosorok, M. R. (2015). New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association*, 110, 583–598.